arXiv:1405.3034v2 [stat.CO] 14 May 2014

G-AMA: Sparse Gaussian graphical model estimation via alternating minimization

Onkar Anant Dalal
Institute for Computational & Mathematical Engg,
Stanford University, USA
onkar@alumni.stanford.edu
Bala Rajaratnam
Department of Statistics,
Stanford University, USA
brajarat@stanford.edu

Abstract

Several methods have been recently proposed for estimating sparse Gaussian graphical models using ℓ_1 regularization on the inverse covariance matrix. Despite recent advances, contemporary applications require methods that are even faster in order to handle ill-conditioned high dimensional modern day datasets. In this paper, we propose a new method, G-AMA, to solve the sparse inverse covariance estimation problem using Alternating Minimization Algorithm (AMA), that effectively works as a proximal gradient algorithm on the dual problem. Our approach has several novel advantages over existing methods. First, we demonstrate that G-AMA is faster than the previous best algorithms by many orders of magnitude and is thus an ideal approach for modern high throughput applications. Second, global linear convergence of G-AMA is demonstrated rigorously, underscoring its good theoretical properties. Third, the dual algorithm operates on the covariance matrix, and thus easily facilitates incorporating additional constraints on pairwise/marginal relationships between feature pairs based on domain specific knowledge. Over and above estimating a sparse inverse covariance matrix, we also illustrate how to (1) incorporate constraints on the (bivariate) correlations and, (2) incorporate equality (equisparsity) or linear constraints between individual inverse covariance elements. Fourth, we also show that G-AMA is better adept at handling extremely ill-conditioned problems, as is often the case with real data. The methodology is demonstrated on both simulated and real datasets to illustrate its superior performance over recently proposed methods.

1 Introduction

In this paper, we consider the problem of sparse inverse covariance estimation for undirected Gaussian graphical models using ℓ_1 regularized maximum likelihood estimation (MLE). Given n realizations of a p dimensional Gaussian random vector, with population and sample covariance matrix denoted by S and Σ respectively, the goal is to estimate $\Sigma \in \boldsymbol{S}^p_{++}$ so that $\hat{\Sigma}^{-1}$ is sparse. For the multivariate Gaussian distribution, the sparsity in the inverse covariance is related to the conditional independence among the features. Specifically, two features i and j are conditionally independent, if and only if the corresponding entry in the inverse covariance matrix is zero, i.e.,$\Sigma^{-1}_{ij} = 0$. Since the MLE is formulated in terms of the inverse covariance X, adding an ℓ_1 regularization with a penalty parameter λ induces sparsity in the estimated inverse covariance matrix. The regularized maximum likelihood estimation problem in terms of $\Sigma^{-1} = X$ is given by

$$\begin{aligned} &\underset{X}{\text{minimize}} \;\; -\log\det X + \langle S, X\rangle + \lambda\|X\|_1 \\ &\text{subject to} \;\; X \in \boldsymbol{S}^p_{++} \end{aligned} \tag{1}$$

where $\|X\|_1 = \sum_{i,j}|X_{ij}|$. The optimization problem (1) is a convex minimization problem in X. For $\lambda = 0$, the solution of Problem (1), given by $X^{(0)} = (\hat{\Sigma}^{(0)})^{-1} = S^{-1}$ is an unbiased estimator of the covariance matrix, but is not well defined for $n < p$. However, for any $\lambda > 0$ the estimate is well defined for any n, with sparsity in $\hat{\Sigma}^{-1}$ increasing for higher values of λ.

The dual of problem (1) can be formulated in terms of the covariance matrix $\Sigma = Y$ as

$$\begin{aligned} &\underset{Y}{\text{minimize}} \;\; -\log\det Y - p \\ &\text{subject to} \;\; \|Y - S\|_\infty \leq \lambda. \end{aligned} \tag{2}$$

In this paper, we propose a new algorithm called Graphical-Alternating Minimization Algorithm (G-AMA), which solves (1) using the Alternating Minimization Algorithm (AMA) proposed by [26]. In [26], AMA was shown to solve the dual problem with the Forward-Backward Splitting (FBS) method from [23] and therefore converges linearly for strongly monotone operators. The iterates of G-AMA are always maintain a feasible dual estimate for the covariance matrix $\hat{\Sigma}$. This is specifically useful while solving large problems with budgeted time and early termination. In addition, G-AMA is very fast for very small values of λ where other state-of-the-art methods converges slowly in comparison since the solutions are highly ill-conditioned.

In many modern practical applications, domain specific knowledge on the covariance structure is known but is not always utilized to its full extent. This knowledge is often available on the

covariance matrix itself, as compared to the inverse covariance matrix, since the sample covariance matrix S is still computable even in the sample starved setting albeit being only positive semi-definite. There are many such practical examples. In genomics it is well known that pairwise correlations of gene (expressions) with genes further downstream are smaller in magnitude. In the environmental sciences when modeling spatial fields, it is often the case that correlations between distant spatial points are less pronounced. As a relative by-product of our AMA analysis, we investigate how to modify G-AMA and other methods to incorporate bound constraints on bivariate covariances if such domain specific information is readily available. We demonstrate that the G-AMA formulation facilitates such analysis quite easily.

The remainder of the paper is organized as follows: Section 1 provides a brief survey of prior work for solving the sparse inverse covariance estimation problem. The methodology and details of G-AMA are given in Section 2, followed by connection with other approaches. Section 3 considers convergence analysis of G-AMA and Section 4 extends the results in Section 3 to two generalizations of the main problem. A detailed account of numerical experiments on synthetic as well as real-life datasets is given in Section 5. This section also demonstrates the use of sparse inverse covariance for portfolio optimization. Further details of the G-AMA method, proofs and additional details on the portfolio optimization application are provided in the appendix.

1.1 Prior Work

A number of algorithms have been proposed to solve the primal problem (1) and its dual problem (2). These algorithms can be briefly classified into two classes, namely, the block-coordinate descent methods and the proximal gradient/Newton methods. The proposed method G-AMA belongs to the second class. In this subsection, we briefly describe algorithms that have been proposed thus far.

Block Coordinate Descent Methods

A dual block coordinate descent which solves a box-QP for each coordinate was first proposed in [1]. In [7], Friedman et al. noticed that each step in dual of the box-QP is equivalent to solving a lasso problem, and proposed the glasso algorithm. The block coordinate descent in glasso was improved upon by applying similar techniques to the primal problem as p-glasso and dp-glasso in [19]. These algorithms take a coordinate descent step for each row and corresponding column and iterate cyclically until convergence. These algorithms are shown to converge to the optimal primal or dual variable. However, the rate of convergence of these methods have not been established, thus providing no theoretical guarantees for the speed of convergence.

Proximal Gradient and Proximal Newton-like Methods

In [6], [17], and [18] the authors use Nesterov's smooth approximation methods and its variations to propose algorithms which achieve ϵ-convergence in $\mathcal{O}(1/\epsilon)$ and $\mathcal{O}(1/\sqrt{\epsilon})$ respectively. A proximal method, QUIC, which uses a proximal Newton step was proposed in [11]. The QUIC algorithm is guaranteed to converge Q-quadratically provided the iterates are close enough to the optimal. This algorithm was later generalized in [16]. In recent work [9], a proximal gradient algorithm, G-ISTA, was proposed and shown to achieve global linear convergence. Both QUIC and G-ISTA algorithms have their respective advantages in various sparsity and condition number regimes. Moreover, both G-ISTA and QUIC use the inverse covariance matrix as their operating variable and seem to outperform the coordinate descent glasso methods in numerical examples.

2 G-AMA: An alternating minimization algorithm for sparse graphical model estimation

2.1 Methodology and problem formulation

In this section, we describe our proposed methodology to solve Problem (1) using alternating minimization algorithm [26]. We write problem (1) is composite form using a dummy variable Z, with $h(X) = -\log\det X$, and $g(Z) = \langle S, Z\rangle + \lambda\|Z\|_1$ as

$$\begin{aligned} &\underset{X,Z}{\text{minimize}} \;\; -\log\det X + \langle S, Z\rangle + \lambda\|Z\|_1, \\ &\text{subject to } Z - X = 0. \end{aligned} \tag{3}$$

An AMA iteration updates the primal variables X, Z and the dual variable Y sequentially using

$$\begin{aligned} X_+ &= \underset{X}{\text{argmin}} \;\; -\log\det X - \langle Y, -X\rangle, \\ Z_+ &= \underset{Z}{\text{argmin}} \;\; \langle S, Z\rangle + \lambda\|Z\|_1 - \langle Y, Z\rangle + \frac{\tau}{2}\|X_+ - Z\|_F^2, \\ Y_+ &= Y + \tau(X_+ - Z_+), \end{aligned} \tag{4}$$

where τ is the step-size. A key feature of the optimization problems in the first two steps is that they can be solved analytically to obtain closed form expressions. The optimality conditions for these two problems are

$$0 = -X_+^{-1} + Y,$$

$$0 \in S + \lambda \text{Sign}(Z_+) - Y + \tau(Z_+ - X_+).$$

The second optimality condition can be rewritten using the soft-thresholding function $\mathcal{S}_\lambda(x) = \text{sign}(x)\,(\max\{(|x| - \lambda), 0\})$ as

$$Z_+ = \mathcal{S}_{\lambda/\tau}(X_+ + (Y - S)/\tau),$$

where the soft-thresholding operator is applied entry-wise. In addition, we observe that substituting X_+ and Z_+ in terms of Y yields a one step update for Y_+ which can be written as

$$Y_+ = \min\left\{\max\left\{Y - S + \tau Y^{-1}, -\lambda\right\}, \lambda\right\} + S. \tag{5}$$

We define a clip function as $\mathcal{C}_\lambda(x) = \min\left\{\max\left\{x, -\lambda\right\}, \lambda\right\}$, where the $\min$, $\max$ functions clip x between $-\lambda$ and λ. The clip function is related to the soft-thresholding function via the identity

$$x = \mathcal{S}_\lambda(x) + \mathcal{C}_\lambda(x).$$

This identity is used in the next section for proof of convergence. The details for alternating minimization algorithm for sparse graphical model estimation (G-AMA) are given in Algorithm 1.

Algorithm 1 G-AMA: Graphical - Alternating Minimization Algorithm

input: sample covariance S, regularization λ, tolerance ϵ_{opt}, backtracking constant $c \in (0, 1)$.
initialize: Y_0, $\tau_{0,0} = 1$, $\Delta_{\text{opt}} = 2\epsilon_{\text{opt}}$.
while: $\Delta_{\text{opt}} > \epsilon_{\text{opt}}$,
$X_{k+1} = Y_k^{-1}$,
compute τ_k: Largest feasible step in $\{c^j \tau_{k,0}\}_{j=0,1,\ldots}$ such that $Y_{k+1} \succ 0$ and satisfies (6),
$Z_{k+1} = \mathcal{S}_{\lambda/\tau_k}(X_{k+1} + (Y_k - S)/\tau_k))$,
$Y_{k+1} = Y_k + \tau_k(X_{k+1} - Z_{k+1})$,
$\tau_{k+1,0} = \dfrac{\langle Y_{k+1} - Y_k, Y_{k+1} - Y_k\rangle}{\langle Y_{k+1} - Y_k, X_k - X_{k+1}\rangle}$,
$\Delta_{\text{opt}} = -\log\det Y_{k+1} - p - \log\det X_{k+1} + \langle S, X_{k+1}\rangle + \lambda\|X_{k+1}\|_1$,
endwhile
output: ϵ-optimal solution to problem (1), covariance estimate $\hat{\Sigma}^{(\lambda)} = Y_{k+1}$.

The Algorithm 1 is terminated using an $\epsilon_{\texttt{opt}}$ tolerance on the duality gap. The step size for each iteration is chosen using the Barzilia-Borwein (BB) step (see 29 in Supplemental section) which

is a two point approximation to the secant equation [2]. A backtracking line search is conducted such that the next iterate is positive definite and satisfies a sufficient descent given by

$$-\log\det(Y_{k+1}) \leq -\log\det(Y_k) + \langle Y_{k+1} - Y_k, Y_k^{-1}\rangle + \frac{1}{2\tau}\|Y_{k+1} - Y_k\|_F^2, \tag{6}$$

where the right hand side is a local quadratic approximation of the dual objective around Y_k. In case a number of backtracking steps fail to satisfy these two criteria, a safe step of $\tau = \alpha^2$ is taken based on Lemma 3.4. More details about step size selection are given in the Supplemental section.

For G-AMA, the objective function of the dual problem (2) can be written as $h^*(Y) + g^*(Y)$ with

$$h^*(Y) = -\log\det Y, \quad \text{and} \quad g^*(Y) = \mathbf{1}_{\{\|Y-S\|_\infty \leq \lambda\}}. \tag{7}$$

The gradient and hessian of the smooth dual objective h^* are given by

$$\nabla h^*(Y) = -Y^{-1},$$

and

$$\nabla^2 h^*(Y) = -Y^{-1} \otimes Y^{-1}.$$

In the following section, we show that over any compact subset of $\mathcal{S}^p_{++}$ the function h^* is strongly convex and the gradient ∇h^* is Lipschitz continuous. These properties are useful in establishing global linear convergence of G-AMA using properties of AMA from [26].

3 Convergence analysis and rate of convergence

In this section, we prove theoretical results regarding global convergence of G-AMA. We first show strong convexity and Lipschitz continuity of the gradient of the objective function over any compact domain. Next we show that the iterates of G-AMA belong to a compact domain bounded away from the boundary of the positive definite cone S^p_{++}. Finally we show the linear convergence of the G-AMA iterates. The proofs are given in the Supplementary section.

We begin with the optimality conditions for Problem (1). The subgradient condition gives

$$-X_*^{-1} + S - \lambda \text{Sign}(X_*) \ni 0.$$

This translates to

$$\begin{aligned} X_*(i,j) = 0 &\iff |S(i,j) - X_*^{-1}(i,j)| \leq \lambda \quad \text{and} \\ X_*(i,j) \neq 0 &\iff S(i,j) - X_*^{-1}(i,j) = -\lambda \cdot \text{Sign}\{X_*(i,j)\}. \end{aligned} \tag{8}$$

We now show that the optimal point satisfying the above conditions is a fixed point for the G-AMA iterations and vice versa.

Lemma 3.1. *A matrix X_* is the optimal solution of the Problem (1) satisfying (8) if and only if the inverse $Y_* = X_*^{-1}$ is a fixed point of the G-AMA iteration in (5), i.e.,*

$$Y_* = \mathcal{C}_\lambda \left(\tau Y_*^{-1} + (Y_* - S)\right) + S. \tag{9}$$

Proof. The proof is given in Appendix. □

The existence of fixed point (9) will allow us to exploit arguments similar to those in [9] to prove global linear convergence.

Now, note that the gradient $\nabla \log\det Y = Y^{-1}$ is not Lipschitz continuous over the entire domain of $\boldsymbol{S}_{++}^p$. However, the gradient can be shown to be Lipschitz continuous over the compact domain

$$\mathcal{D} = \{Y \mid \alpha I \preceq Y \preceq \beta I\} \subset \boldsymbol{S}_{++}^p, \quad \text{for} \quad 0 < \alpha < \beta < \infty. \tag{10}$$

Lemma 3.2. *([9, Lemma 2]). For $Y_1, Y_2 \in \boldsymbol{S}_{++}^p$, the gradient $\nabla \log\det Y = Y^{-1}$ satisfies*

$$\frac{1}{\beta^2}\|Y_1 - Y_2\|_2 \leq \|Y_1^{-1} - Y_2^{-1}\|_2 \leq \frac{1}{\alpha^2}\|Y_1 - Y_2\|_2,$$

where $\alpha = \min\{\lambda_{\min}(Y_1), \lambda_{\min}(Y_2)\}$ and $\beta = \max\{\lambda_{\max}(Y_1), \lambda_{\max}(Y_2)\}$.

We also note that the hessian of $\log\det$ function is given by

$$\nabla^2 \log\det Y = -Y^{-1} \otimes Y^{-1}, \tag{11}$$

and the function $\log\det Y$ is strongly convex when Y is restricted to the domain $\mathcal{D}$. Next we show that the iterates of G-AMA belong to the bounded compact set $\mathcal{D}$ defined in (10) and give explicit values for the constants α and β. We begin with the upper bound based on [9, Lemma 8].

Lemma 3.3. *([9, Lemma 8]). Let $0 \prec \alpha_l I \preceq Y_l$ for $l = 0, 1 \ldots k$ and let $\tau_k < \alpha_k^2 < \beta^2$ (for $\beta = \|Y_0 - Y_*\|_F + \|Y_*\|_2$) be the step size used for G-AMA iterations starting with $\alpha_0 I \prec Y_0$. Then*

the next iterates Y_{k+1} satisfies

$$Y_{k+1} \preceq \beta I, \tag{12}$$

and hence by inductive argument $Y_k \preceq \beta I$, for all k.

Proof. The proof is given in Appendix. □

Note that Lemma 3.3 only assumes a strictly positive α_l for $l = 0, 1, \ldots, k$ to prove the upper bound on Y_{k+1}. It does not assume that the subsequent iterates are strictly bounded away from the semidefinite boundary. This result is shown in the next lemma which will use a universal upper bound β on $\|Y_k\|_2$, for all $(k = 0, 1, \ldots)$. Based on Lemma 3 in [11] we show that the iterates of G-AMA are bounded away from the boundary of the positive definite cone and satisfy $\alpha I \preceq Y_k$ for some universal $\alpha > 0$ for all $(k = 0, 1, \ldots)$.

Lemma 3.4. *([11, Lemma 3]). The set $\mathcal{U} = \{Y \in \boldsymbol{S}^p_{++} \mid Y \prec \beta I$ and $\log\det Y_0 < \log\det Y\}$ satisfies $\alpha I \preceq Y$ with $\alpha = \frac{\lambda^p}{\beta^{(p-1)}}$, where β is the upper bound from (12) of Lemma 3.3.*

Proof. The proof is given in Appendix. □

We now show global linear convergence of G-AMA using Lemma 3.3, and Lemma 3.4.

Theorem 3.5. *The iterates Y_k of Algorithm 1 satisfy $\alpha I \preceq Y_k \preceq \beta I$ as well as*

$$\|Y_{k+1} - Y_*\|_F \leq \max\left\{\left|1 - \frac{\tau_k}{\alpha^2}\right|, \left|1 - \frac{\tau_k}{\beta^2}\right|\right\} \|Y_k - Y_*\|_F, \tag{13}$$

for $k = 0, 1 \ldots$. More specifically, for step size $\tau_k < \alpha^2$,

$$\|Y_{k+1} - Y_*\|_F \leq \gamma \|Y_k - Y_*\|_F, \tag{14}$$

where $\gamma < 1$ and hence the iterates converges to a ϵ-optimal solution Y_* in $\mathcal{O}\left(\log(1/\epsilon)\right)$ iterations.

Proof. The proof is given in Appendix. □

The maximum step size α^2 provided by α from Lemma 3.4 is very conservative. We show later that much better performance can be achieved in practice using heuristics for choosing the step size.

3.1 Connections to G-ISTA

The dual objective functions h^* and g^* from (7) are conjugate functions (see [4]) of the primal objective functions

$$h(X) = -\log\det X, \quad \text{and} \quad g(Z) = \langle S, Z\rangle + \|Z\|_1.$$

In [9], Problem (1) is solved using a proximal gradient method (FBS also known as ISTA) and global linear convergence of the algorithm is established using strong convexity of h and Lipschitz continuity of the gradient, ∇h, over a compact domain. Given an initial point, the subsequent iterates are shown to remain in a compact domain and hence the proximal gradient method converges linearly [22]. The two algorithms G-AMA and G-ISTA are closely related, and as expected there are some parallels in their respective convergence analysis. The parallels between the two methods can also be used to choose the step size using an equivalent Barzilia-Borwein step.

3.2 Connections to QUIC and DC-QUIC

An extension wrapper for QUIC called DC-QUIC was proposed in [10] that provides a divide-and-conquer framework which approximates a large inverse covariance estimation problem with multiple small problems. These problems are then solved using standard QUIC and the solutions are concatenated to get an approximate solution of the original problem. More specifically, the DC wrapper is an independent top layer that only uses QUIC as its base algorithm. Any equivalent method like G-AMA can replace QUIC in this bigger divide-and-conquer framework. In this light, our current work compares the base algorithms of G-AMA and QUIC, and the speedup of G-AMA over QUIC can therefore be also directly utilized in the spirit of a "DC-GAMA".

3.3 Connections to ALM [24]

We also note that the G-AMA algorithm is different from the ALM method [24] in several fundamental aspects, including (a) Theoretical, (b) Algorithmic and, (c) Numerical. In particular, (a) The ALM algorithm is shown to attain an ϵ-optimal solution in O($1/\epsilon$) iterations. In comparison, G-AMA is shown to attain an ϵ-optimal solution in O($\log(1/\epsilon)$) iterations. The rate of convergence of our proposed method G-AMA is therefore faster, and it is not faster just by an order of magnitude (such as O($1/\epsilon^2$) of Nesterov's accelarated gradient methods), it is in a different regime altogether. In particular, G-AMA has linear convergence. (b) The ALM algorithm is an implementation of ADMM [3], whereas G-AMA is an implementation of AMA [26]. The key difference between these methods is the first subproblem in (4) where the ADMM algorithm uses a quadratic term

and AMA does not. As we shall see this very vital difference provides the speedup enjoyed by G-AMA as a result of the strongly convex objective function. (c) Numerically, ALM was substantially shown to be slower than QUIC in [11] and was therefore not considered state-of-the-art in comparisons with G-AMA. To be fair, G-AMA will be compared to the more superior QUIC and G-ISTA algorithms and will be shown to outperform these current state-of-the-art methods.

3.4 Choice of penalty parameter and Statistical Consistency

Once convergence of G-AMA to a global minimum is established, the next natural question to investigate is the choice of penalty parameter λ. There are various approaches that have been proposed in this regard. First, cross validation is a popular technique that is often used. Second, one could use a value for the penalty parameter λ that yields a desired level of sparsity (such as 3% edge density). Third, Bayesian information type criteria are also popular for selection of the appropriate level of regularization [12]. Fourth, the penalty parameter can be chosen so as to control the probability of false detection of edges, i.e., a certain predetermined level of error control [1]. An important question pertains to asymptotic recovery of the underlying graph when the data is generated from a sparse Gaussian graphical model (i.e., model selection consistency). Since G-AMA yields the ℓ_1 regularized MLE (similar to G-lasso, G-ISTA and other approaches), we can leverage already established statistical theory to assert asymptotic consistency of G-AMA estimates. The reader is referred to [1] in this regard, who show that provided the dimension $p = \mathcal{O}(n^\gamma)$ for some $\gamma > 0$, the ℓ_1 regularized MLE recovers the underlying graph with probability tending to 1.

4 Generalized G-AMA

We now proceed to demonstrate that the G-AMA framework is even richer than the formulation in Problem (3). In this section, we extend our G-AMA algorithm to solve two natural generalizations of the sparse covariance estimation problem.

The first generalization modifies the dual problem (2) to include explicit constraints on the bivariate covariances in Σ. As discussed in Section 1, scenarios where domain specific knowledge regarding correlation structure is available arise in numerous applications. This knowledge/structure is more easily found in bivariate pairwise/marginal relationships between feature pairs since such quantities can be simply calculated from the sample covariance matrix, as compared to the sample inverse covariance matrix which is not defined when $n < p$. Since the operating variable for G-AMA is the estimate of the covariance matrix $\hat{\Sigma} = Y$, it can easily incorporate this domain

specific knowledge by adding constraints on the covariance matrix (compared to the inverse covariance matrix). Thus the G-AMA framework allows for an additional layer of regularization, while maintaining the sparsity in the inverse covariance matrix, and in the process enriches the type of graphical modeling that can be achieved. We develop a comprehensive methodology in this regard that extends the standard inverse covariance estimation problem. The dual method for glasso can also be modified to incorporate simple bound constraints like $l_{ij} \leq \sigma_{ij} \leq u_{ij}$. However, a more complicated constraint would lose the fast coordinatewise descent algorithm for the lasso subproblem.

The second generalization modifies the primal problem (1) with a ℓ_1 regularization on a symmetric linear transform of the inverse covariance matrix X (see Problem (19)). In many applications the underlying graph structure is known to be symmetric, i.e., the edge weights (partial covariance/correlations) assume equal values (equisparsity) either naturally or by design (see [8] and references therein for examples). In such settings, it is often both necessary and useful to estimate a partial correlation graph that respects these symmetries (equisparisty). Such additional constraints on graphical models are both of theoretical and applied interest due to the added layer of regularization and structure it provides. This type of information about the structure of the underlying graph can be incorporated by using a regularization term in the form of a symmetric linear transform.

4.1 Methodology

Generalization of convex constraint on Σ

The constraint set on Σ in Problem (2) can be generalized to any convex constraint $\mathcal{D}$. In this subsection we will illustrate the extension of G-AMA for a problem with arbitrary bound constraints on the bivariate covariances. We also show that constraints on the covariance matrix can sometimes be translated directly to constraints on the inverse covariance, hence primal methods can also be somewhat useful in this context. The dual problem can be modified with generalized bound constraints as

$$\begin{aligned} &\underset{Y}{\text{minimize}} \;\; -\log\det Y \\ &\text{subject to} \;\; l_{ij} \leq Y(i,j) \leq u_{ij}. \end{aligned} \tag{15}$$

Here, each bivariate covariance σ_{ij} is allowed to vary in the interval $[l_{ij}, u_{ij}]$ which is a generalization of the interval $[s_{ij} - \lambda, s_{ij} + \lambda]$ from the dual of regularized MLE problem (2). These modified constraints break away from the maximum likelihood framework and allows a different

kind of regularization based on domain specific knowledge. The corresponding primal problem formulated in terms of the inverse covariance matrix is given by

$$\begin{aligned} &\underset{X}{\text{minimize}} \ -\log\det X + \langle \bar{S}, X\rangle + \sum_{i,j} \bar{\lambda}_{ij}|X_{ij}| \\ &\text{subject to} \quad X \in \boldsymbol{S}^p_{++} \end{aligned} \tag{16}$$

where $\bar{S}$ is a modified sample covariance (midpoints of the intervals) and the third term is a ℓ_1 penalty with more flexible regularization $\bar{\lambda}_{ij}$ (half length of the intervals) for X_{ij}. These are related to the lower and upper bounds of the dual formulation as

$$\bar{S}_{ij} = \frac{l_{ij} + u_{ij}}{2}, \ \text{ and } \ \bar{\lambda}_{ij} = \frac{u_{ij} - l_{ij}}{2}. \tag{17}$$

Note that unlike the MLE formulation, where $S \succeq 0$ provided a feasible point $S + \lambda I \succ 0$ (for any $\lambda > 0$), the generalized problem is no longer guaranteed to be feasible for arbitrary choice of parameters l_{ij}'s and u_{ij}'s. For this problem, the algorithm update for Y is given by

$$Y_+ = \Pi_{\mathcal{D}} \left\{ Y + \tau Y^{-1} \right\}, \tag{18}$$

where $\Pi_{\mathcal{D}}$ is the projection onto the convex constraint set $\mathcal{D} = \{x \mid l_{ij} \leq x_{ij} \leq u_{ij}, \forall i, j\}$ defined by the bound constraints. The simplified projection for the bound constraints defined in (15) is given by

$$\{Y_+\}_{ij} = \min\left\{ \max\left\{ \left(Y + \tau Y^{-1}\right)_{ij}, l_{ij} \right\}, u_{ij} \right\}.$$

The methodology for choosing the step size remains the same as for the standard G-AMA algorithm and the linear convergence results hold.

It is important to note that the generalization to any convex constraint set $\mathcal{D}$, in addition to the bound constraints or the standard constraint set of $\mathbf{1}_{\{\|Y-S\|_\infty \leq \lambda\}}$, might not always be easily translated to the primal problem. This fact underscores the flexibility of the above generalized G-AMA formulation on the dual.

Generalization of ℓ_1 penalty on Ω

Consider a generalized regularization on the ℓ_1 norm of A^*XB instead of the ℓ_1 norm of X. Then a generalized version of problem (1) is given by

$$\begin{aligned} &\underset{X}{\text{minimize}} \;\; -\log\det X + \langle S, X\rangle + \lambda\|A^*XB\|_1 \\ &\text{subject to} \;\; X \succ 0. \end{aligned} \tag{19}$$

It can be easily shown that the above formulation enforeces equality (equisparsity or linear constraints) on the elements of the inverse covariance matrix. As in (3), the problem can be decomposed into composite form using the constraint $A^*XB - Z = 0$. Applying the alternating minimization method, yields the following update steps

$$\begin{aligned} X_+ &= \underset{X}{\text{argmin}} \;\; -\log\det X + \langle S, Z\rangle - \langle Y, -A^*XB)\rangle, \\ Z_+ &= \underset{Z}{\text{argmin}} \;\; \lambda\|Z\|_1 - \langle Y, Z\rangle + \frac{\tau}{2}\|Z - A^*X_+B\|_F^2, \\ Y_+ &= Y + \tau\left(A^*X_+B - Z_+\right). \end{aligned} \tag{20}$$

The optimality conditions for the first two steps are given by

$$\begin{aligned} 0 &= -X_+^{-1} + S + BYA^*, \\ 0 &\in \lambda\text{sign}(Z_+) - Y + \tau\left(Z_+ - A^*X_+B\right). \end{aligned} \tag{21}$$

The second optimality condition can be rewritten using the soft-threshold function as

$$Z_+ = \mathcal{S}_{\lambda/\tau}\left(A^*X_+B + Y/\tau\right). \tag{22}$$

Substituting the two optimality conditions for X_+ and Z_+ gives a direct update for Y_+ in terms of Y as

$$\begin{aligned} Y_+ &= Y + \tau\left[A^*X_+B - \mathcal{S}_{\lambda/\tau}\left(A^*X_+B + Y/\tau\right)\right] \\ &= \mathcal{C}_\lambda\left[Y + \tau A^*(S + BYA^*)^{-1}B\right]. \end{aligned} \tag{23}$$

4.2 Convergence Analysis

The proof of convergence from the previous section can be generalized to apply to the modified G-AMA described in Section 4.1 and 4.1. Since the new iterates belong to a subset of the set $\mathcal{U}$

defined in Lemma 3.4, the same lemma is valid for the new problems. Lemma 3.2, Lemma 3.3 also hold for the modified method. In order to prove Lemma B.1, we first recall the non-expansiveness of the prox operator for any convex set $\bar{\mathcal{D}}$ [22].

Lemma 4.1. *([22]): The proximal operator $prox_{\tau\bar{\mathcal{D}}}$ for an indicator function of a convex set $\bar{\mathcal{D}}$ given by*

$$prox_{\tau\bar{\mathcal{D}}}(x) = \underset{y}{argmin} \left\{ \mathbf{1}_{\bar{\mathcal{D}}} + \frac{\tau}{2}\|y - x\|^2 \right\} = \Pi_{\bar{\mathcal{D}}}(x) \tag{24}$$

satisfies

$$\|prox_{\tau\bar{\mathcal{D}}}(x) - prox_{\tau\bar{\mathcal{D}}}(y)\|_2 \leq \|x - y\|_2 \tag{25}$$

Lemma 4.1 is useful for proving linear convergence for a generalized convex constraint on Σ. We now formally show global linear convergence of generalized G-AMA.

Theorem 4.2. *The new iterate Y_+ of the generalized G-AMA algorithm from (18) and (23) satisfy*

$$\|Y_+ - Y_*\|_F \leq \gamma \|Y - Y_*\|_F, \tag{26}$$

for some $\gamma < 1$ and hence the iterates converges to a ϵ-optimal solution Y_* in $\mathcal{O}\left(\log(1/\epsilon)\right)$ iterations.

Proof. The proof is given in the Appendix. □

5 Numerical Experiments

The performance of G-AMA was compared to other algorithms on synthetic as well as real data sets. We begin with illustrating linear convergence of G-AMA and how the convergence of G-AMA varies with condition number of the estimate $\hat{\Sigma}^*$. The left hand plot in Fig. 1 illustrates the convergence of $\|Y_k - Y_*\|_F$ and the right hand plot demonstrates the convergence of the G-AMA iterates with respect to the duality gap Δ_{opt}. The speed of the algorithm becomes slower as the condition number $\kappa(\Sigma^*)$ increases.

Next, we demonstrate timing comparisons between G-AMA, QUIC and G-ISTA for solving Problem (1). We test the three algorithms on synthetic as well as real data sets. The timings reported in this section are the wall-clock timings and not the CPU timings. The three algorithms use LAPACK library for computing Cholesky factors and inversion which leverages multiple processors when available. This benefits the first order methods, G-AMA and G-ISTA which expend

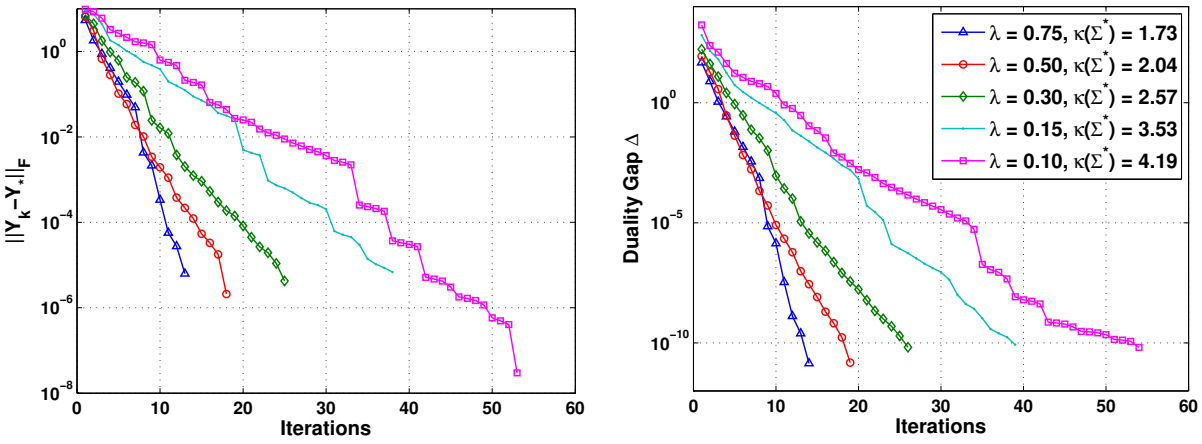

Figure 1: Convergence of G-AMA: (a) $\|Y_k - Y_*\|_F$ and (b) Duality gap Δ_{opt}

majority of their computational resources in these operations, in contrast to QUIC which expends computational resources in solving the lasso subproblem.

5.1 Synthetic datasets

The first set of comparison experiments were conducted on synthetic data. For this subsection, the data was generated using the procedure in [18]. For a given problem size p, the underlying inverse covariance matrix $\Omega = \Sigma^{-1}$ was generated by choosing the off-diagonal entries i.i.d. from a uniform distribution over $[-1, 1]$. To obtain a desired sparsity level of $sp\%$ (percentage of non-zero entries in $\hat{X}$), the off-diagonals were set to zero with probability $sp/100$ (As in [9], sparsity levels of 3% and 15% were used). A multiple of the identity matrix was added to this matrix to adjust the smallest eigenvalue of $\Omega = \Sigma^{-1}$ to 1. This procedure ensures that Ω is positive definite, sparse, and well-conditioned. Finally, datasets of n i.i.d. samples were generated from the normal distribution $\mathcal{N}_p(0, \Omega^{-1})$. For each Ω, sample sizes of $n = 0.2p$ and $n = 1.2p$ were chosen to illustrate both when $n << p$ and when $n \geq p$.

Heuristics and insights

During the numerical experiments it was observed that as λ increases, the number of iterations go down. Moreover, for large values of λ, the performance of each algorithm depends on the initial point that is chosen. While G-ISTA and QUIC use $\Omega_0 = (\texttt{diag}\,(S)+\lambda I)^{-1}$ as the starting point, the inverse $Y_{0,1} = (\texttt{diag}\,(S) + \lambda I)$ works very well for G-AMA for large values of λ. However, as the value of λ decreases, it was observed that the next iterate Y_1 was not always positive definite.

This can be explained by violation of dual feasibility of Y_0 for small values of λ. A better consistent starting point which works for all values of λ is $Y_{0,2} = S + \lambda I$. The timing comparisons shown in this section display the time required for convergence starting with the best initial point. In the majority of cases, the best initial point was $S + \lambda I$ for G-AMA. As a general guideline, $Y_{0,2}$ can be used in general, and $Y_{0,1}$ can also be used to obtain some speedup when the expected solution is very sparse ($< 2\%$). We also observe that as p increases G-AMA and G-ISTA, scale more favorably than QUIC. A general observation as n increases is that all the methods perform better. This can be explained by the fact that condition number of the sample covariance S improves with n. Therefore, the iterates and the optimal solution have better conditioning.

G-AMA vs G-ISTA vs QUIC

Similar to G-ISTA and QUIC, the G-AMA was implemented in C with a wrapper for *MATLAB* and the three methods were compared in terms of time taken on a machine running the 32-bit version of Ubuntu 12.10 with an Intel Core i7 870 CPU and 8 GB RAM. Table 1 and Table 2 display the results for problems with $p = 2000$ and $p = 5000$ for varying n and sp.

Table 1: Time (Iterations) and Duality Gap Comparison for G-AMA with G-ISTA and QUIC

		algorithm:		G-ISTA		QUIC		G-AMA	
problem	λ	nnz(Z_*)	$\kappa(Y_*)$	**time/iter**	**gap**	**time/iter**	**gap**	**time/iter**	**gap**
$p = 2000$; $n = 400$; $sp = 03\%$	0.12	2.36%	2.24	17s/27	5e-10	28s/9	8e-10	**6s**/15	7e-11
	0.09	7.24%	6.92	56s/75	1e-10	99s/10	1e-10	**27s**/57	6e-11
	0.06	15.15%	19.13	51s/98	3e-07	331s/13	3e-07	**28s**/59	5e-09
	0.03	27.76%	45.67	327s/457	4e-06	1853s/24	4e-06	**18s**/43	3e-06
$p = 2000$; $n = 2400$; $sp = 03\%$	0.12	0.10%	1.14	9s/18	1e-12	**3s**/6	1e-12	**3s**/6	1e-12
	0.09	0.91%	1.48	21s/31	2e-12	9s/7	2e-12	**6s**/13	1e-12
	0.06	3.06%	2.76	22s/30	1e-09	33s/8	2e-08	**18s**/32	1e-08
	0.03	14.56%	9.87	37s/53	2e-07	585s/16	3e-07	**25s**/41	4e-08
$p = 2000$; $n = 400$; $sp = 15\%$	0.12	2.83%	2.80	8s/15	1e-06	24s/7	5e-06	**7s**/16	2e-07
	0.09	7.80%	11.67	28s/49	4e-08	112s/10	1e-07	**22s**/49	2e-08
	0.06	15.23%	28.07	90s/150	8e-07	387s/13	1e-06	**37s**/63	4e-07
	0.03	27.37%	63.37	654s/813	1e-06	1990s/24	1e-06	**41s**/67	4e-08
$p = 2000$; $n = 2400$; $sp = 15\%$	0.12	0.03%	1.10	3s/8	1e-12	3s/6	1e-12	**2s**/4	1e-12
	0.09	0.63%	1.34	6s/14	1e-09	11s/9	3e-09	**3s**/7	1e-09
	0.06	5.38%	4.01	26s/32	9e-08	129s/12	1e-07	**16s**/31	5e-09
	0.03	19.92%	17.55	87s/113	5e-07	351s/11	3e-06	**19s**/36	1e-06

The first order methods G-AMA and G-ISTA are terminated using a tolerance condition on the duality gap. However, the termination of QUIC is controlled by the tolerance on sub-gradient condition only. The problems were solved using QUIC by setting the tolerance $\epsilon_{\text{tol}} = 10^{-10}$ and the corresponding duality gap achieved by QUIC was then used for G-AMA, and G-ISTA. The

Table 2: Time (Iterations) and Duality Gap Comparison for G-AMA with G-ISTA and QUIC

		algorithm:		G-ISTA		QUIC		G-AMA	
problem	λ	nnz(Z_*)	$\kappa(Y_*)$	**time/iter**	**gap**	**time/iter**	**gap**	**time/iter**	**gap**
$p = 5000$; $n = 1000$; $sp = 3\%$	0.10	0.47%	1.43	59s/11	2e-07	102s/7	3e-07	**41s**/7	7e-08
	0.08	2.03%	2.29	215s/28	1e-11	434s/9	1e-11	**92s**/16	1e-11
	0.06	6.36%	8.29	478s/54	9e-09	1607s/10	9e-09	**284s**/47	7e-09
	0.04	13.69%	23.12	935s/127	6e-07	6315s/14	1e-06	**322s**/53	1e-06
$p = 5000$; $n = 6000$; $sp = 3\%$	0.08	0.16%	1.15	76s/11	2e-08	40s/5	9e-08	**31s**/5	1e-09
	0.06	1.12%	1.59	59s/11	2e-07	145s/6	1e-06	**52s**/9	9e-09
	0.04	3.24%	3.65	213s/27	2e-09	563s/8	3e-09	**167s**/27	8e-10
	0.02	12.89%	15.03	391s/61	4e-06	4682s/12	4e-06	**247s**/41	1e-06
$p = 5000$; $n = 1000$; $sp = 15\%$	0.08	2.53%	3.31	**84s**/12	6e-07	365s/7	9e-07	135s/20	7e-07
	0.06	6.94%	18.20	933s/109	1e-08	2099s/11	1e-08	**507s**/82	7e-09
	0.04	13.88%	39.37	1599s/229	4e-05	5270s/12	4e-05	**270s**/45	4e-05
	0.02	26.02%	84.12	6865s/1000	4e-03	31841s/24	3e-06	**390s**/62	2e-07
$p = 5000$; $n = 6000$; $sp = 15\%$	0.08	0.04%	1.10	66s/11	1e-10	31s/5	3e-09	**24s**/4	4e-11
	0.06	0.72%	1.43	53s/10	8e-07	79s/5	1e-05	**35s**/6	8e-08
	0.04	5.50%	7.14	236s/31	1e-07	1386s/10	8e-07	**218s**/37	9e-07
	0.02	18.66%	26.96	1771s/218	4e-07	7132s/12	4e-06	**225s**/39	3e-07

times reported for G-AMA and G-ISTA in Tables 1 and 2 are the times required to achieve a better duality gap than QUIC. This explains the different duality gaps reported for each of the three methods.

Note that the time required for convergence using G-AMA is substantially lower than G-ISTA and QUIC in all instances except one. Moreover, the advantage of G-AMA is even more prominent when λ is small and/or when the solution is ill-conditioned. As discussed in the next Subsection, this benefit of G-AMA is magnified when working with real data where the optimal solutions are often highly ill-conditioned.

5.2 Real dataset: *Estrogen* and *Temperature*

We compare the three methods on two real datasets *Estrogen* from [21] and *Temperature* from [5]. The *Estrogen* dataset consists of expression data for $p = 682$ genes from $n = 158$ patients with breast cancer. The *Temperature* dataset consists of average annual temperature measurements from $p = 1732$ locations over $n = 157$ years (1850-2006). The values of λ chosen to test the algorithms varied in order to get sparsity between $2\% - 12\%$. Both datasets yield extremely ill-conditioned problems for the above values of λ. We know that the ability of G-AMA to control the duality gap is extremely useful for attaining accurate solutions in such highly unstable regime. A maximum of 5000 iterations were used for G-ISTA and G-AMA.

It was observed that the duality gap achieved by QUIC decreases with tolerance on the sub-

Table 3: Wall times, iterations and duality gap comparison for G-AMA with G-ISTA and QUIC. *(Note that the reported duality gap for QUIC cannot be reduced further for any tolerance set for the solver. For timing comparison purposes, G-AMA therefore reports two times. The first achieves a duality gap at least as small as QUIC. The last column* reports the time to achieve a significantly lower fixed duality gap of $\epsilon_{opt} = 10^{-10}$, a gap that QUIC cannot achieve for any subgradient tolerance. G-ISTA is significantly slower or fails to converge within 5000 iterations for these cases.)*

	algorithm:		G-ISTA		QUIC		G-AMA		
$\boldsymbol{\lambda}$	nnz(Z_*)	$\kappa(Y_*)$	**time/iter**	**gap**	**time/iter**	**gap**	**time/iter**	**gap**	**time**
Dataset: Estrogen p = 682; n = 158									
0.40	2.62%	42.36	21s/532	1e-06	**3s**/13	2e-06	5s/180	2e-06	**10s***
0.30	3.39%	88.57	35s/911	7e-07	**6s**/19	2e-07	17s/548	4e-08	**21s***
0.20	4.39%	192.70	80s/2182	7e-07	23s/27	8e-07	**16s**/488	8e-07	**27s***
0.10	7.10%	475.24	NA/5000	NA	64s/49	2e-06	**11s**/339	2e-06	**22s***
0.05	12.36%	985.99	NA/5000	NA	335s/93	2e-06	**11s**/321	1e-06	**16s***
Dataset: Temperature p = 1732; n = 157									
0.30	1.75%	589.70	1242s/2940	5e-07	**279s**/21	1e-06	438s/1161	9e-07	**639s***
0.20	2.02%	1076.07	NA/5000	NA	589s/28	4e-06	**356s**/931	3e-06	**642s***
0.10	3.30%	2301.98	NA/5000	NA	953s/47	2e-05	**216s**/581	1e-05	**301s***
0.05	5.59%	4505.01	NA/5000	NA	3402s/90	8e-06	**249s**/652	6e-06	**374s***
0.03	7.84%	7201.09	NA/5000	NA	9433s/136	2e-05	**209s**/580	9e-06	**346s***

gradient. However, it stalls at a particular point despite decreasing the tolerance provided to the solver. The duality gaps reported in Table 3 for QUIC are the best that could be achieved by any tolerance on subgradient condition. These are as high as $2 \times 10^{-6} - 2 \times 10^{-5}$ when the problem is highly ill-conditioned. In contrast, there are two timings reported for G-AMA. Similar to synthetic experiments, the first is the time required to achieve a duality gap at least as small as the gap achieved by QUIC. The second time (last column*) is that required to achieve a fixed duality gap of $\epsilon = 10^{-10}$. We note that QUIC is unable to achieve this level of duality gap for any subgradient tolerance, whereas G-ISTA is very slow in such problems. G-AMA suffers from neither of these two issues.

Now, as seen in Table 3, for both examples the condition number of the optimal solution is extremely high for lower values of $\lambda = 0.20, 0.10, 0.05, 0.03$. The QUIC algorithm performs better than G-AMA when λ is higher, but is extremely slower in comparison as λ reduces. For the *Estrogen* dataset, G-AMA is up to 30 times faster than QUIC when the condition number of solution is almost 1000. For the *Temperature* dataset, when the sparsity level is around 8% for $\lambda = 0.03$, G-AMA is 45 times faster than QUIC. Such speedup is extremely vital for applications where regularized maximum likelihood estimation problem need to be solved several times. In particular, determining the level of the penalty parameters is often done through cross-validation which in turn requires solving the ℓ_1 regularized maximum-likelihood estimation problem for a

grid of values of the penalty parameter [see for example [25] for the climate field reconstruction application]. In addition, when uncertainty quantification is required, both parametric and non-parametric bootstrap require solutions of the optimization problem numerous times.

5.3 Portfolio optimization using sparse inverse covariance estimates

We now apply G-AMA for solving the regularized maximum likelihood covariance estimation problem in the context of financial portfolio optimization. A regularized inverse covariance estimate is a critical ingredient in determining the minimum variance portfolio. Moreover, extremely fast estimation methods like G-AMA are required for such problems as covariance matrices need to be estimated repeatedly over moving time blocks. In this section we shall compare the performance of G-AMAwith popular methods in the literature to illustrate its efficacy in the financial portfolio optimization context.

Minimum variance portfolio with rebalancing

A portfolio optimization problem refers to the problem of determining weights or proportions (in monetary terms) in order to invest in a set of securities that minimize the risk for a given level of return. We will focus on the minimum variance portfolio using the covariance estimate computed by solving problem 1. For a portfolio with p risky assets, let r_i denote the return of asset i over a given period, *i.e.*, its change in price over one time period divided by its price at the beginning of the period. Let Σ denote the covariance matrix of $r = (r_1, r_2, \ldots, r_p)$ and w_i be the weight of asset i in the portfolio during a given period (w_i could be positive or negative based on long or short positions). The minimum variance portfolio selection problem solves

$$\begin{aligned} &\underset{w}{\text{minimize}}\ w^T \Sigma w \\ &\text{subject to } \mathbf{1}^T w = 1, \end{aligned} \tag{27}$$

where the objective denotes the variance of the return (also defined as the risk) associated with a particular portfolio. The linear constraint represents the budget constraint. For a given Σ, this quadratic program has a closed form analytic solution $w^* = (\mathbf{1}^T \Sigma^{-1} \mathbf{1})^{-1} \Sigma^{-1} \mathbf{1}$.

The standard problem defined above assumes stationarity of the returns. In order to account for non-stationarity, we employ the minimum variance portfolio rebalancing (MVR) strategy. This approach updates the portfolio weights every L units of time, effectively dividing the trading horizon into blocks each consisting of L time units. At the start of each block, the minimum variance portfolio problem is solved based on the past estimation horizon size of N_{estim} observations of

returns. The weights w are then held constant for L time units during the holding periods. We assume that the total number of time units in the entire period, $N_{\text{tot}} = N_{\text{estim}} + K \cdot L$, for some positive integer K. Therefore, we have K updates of portfolio weights, say $w^{(j)}$ over the holding period of $[N_{\text{estim}} + 1 + (j-1) \cdot L, N_{\text{estim}} + j \cdot L]$ for $j = 1, \ldots K$. Here, $w^{(j)}$ is calculated using the estimated covariance of asset returns over the jth holding period.

Application to the DJIA

In the numerical study, we use the stock data for 29 stocks that constituted the Dow Jones Industrial Average (DJIA) as of July 2008, and had a start date of January 01, 1995 and an end date of October 26, 2012 (4473 trading days). We have chosen this basket of securities deliberately for comparison purposes as covariance estimation methods have been recently illustrated on this data. The first rebalancing interval begins from January 01, 1995 and the last interval begins on July 02, 2012. We have,

$$N_{\text{tot}} = 4473, \quad K = 58, \quad L = 80, \quad N_{\text{estim}} = 75, 150, 225, 300.$$

We estimate the sample covariance S using N_{estim} weekly returns from the DJIA. The entire trading horizon consists of $K = 58$ holding periods of $L = 80$ days (16 weeks) each. We compare the MVR strategy using covariance matrices generated by SparseConc using G-AMA (with constant penalty parameter $\lambda = 0.1\|S\|_2$) versus covariance estimation methods considered in [27]. These include **CondReg** (condition number regularized covariance from [20]), **Factor** (a shrinkage method from [13]), **LedoitWolf** (a linear shrinkage scheme from [14]), **Sample** (sample covariance matrix) as well as the **DJIA**. The performance of these methods is compared on the metrics of *Realized Return*, *Realized Risk*, *Realized Sharpe Ratio* (ratio of excess expected return and risk), *Turnover* (measure of portfolio switching for each time-interval), *Normalized Wealth Growth* and the *Size of the Short Side* (measure of negative weights chosen). The formal definitions for the performance metrics are given in Appendix C.

Performance metrics

Fig. 2 shows the normalized wealth growth over the trading horizon for five methods and the DJIA for four different values of N_{estim}. As can be seen, the wealth growth using G-AMA estimates give the best results compared to other methods. Table 4 shows the values for realized return, realized risk, Sharpe ratio, turnover and size of short side. Each entry is the mean (standard deviation) of the corresponding metric over the trading period. The standard deviations are computed in a heteroskedasticity-and-autocorrelation consistent manner discussed in [15, Sec. 3.1]. The Turnover values are illustrated in Fig. 3.

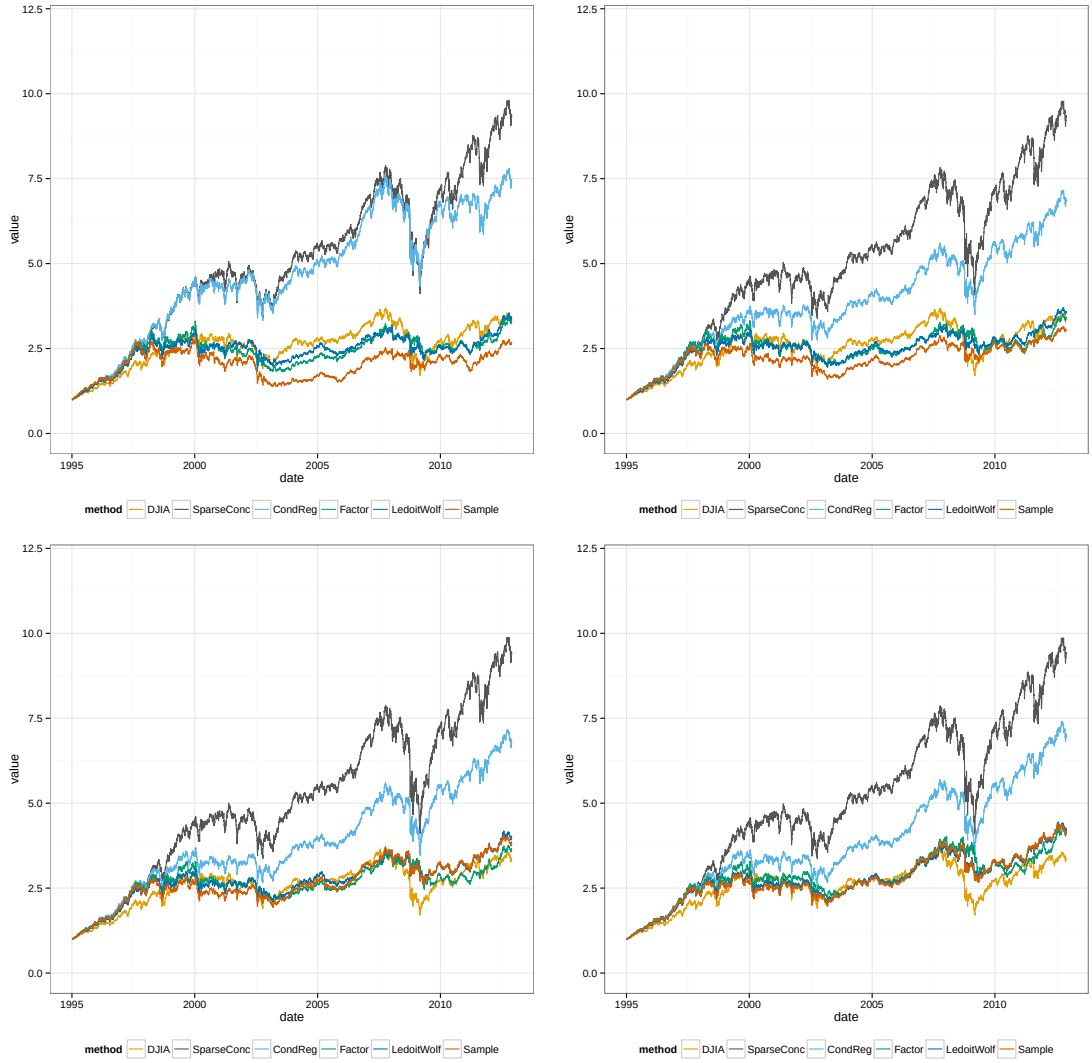

Figure 2: Normalized wealth growth of G-AMA vs other methods ($\lambda = 0.10$)

Comparison Results

The normalized wealth growth for G-AMA is substantially better than other methods. G-AMA also gives high values of Sharpe ratio across estimation horizons. An advantage of G-AMA is that the turnover is lower than for other methods and the size of short-side is negligible. G-AMA gives rise to a stable portfolio which avoids excessive continuous re-balancing, transaction and borrowing costs. So in this sense, G-AMA's net effective wealth growth are even higher than those depicted in Fig. 2.

Table 4: Realized Return, Realized Risk, Sharpe Ratio, Turnover and Size of Short Side

covariance regularization scheme	return [%]	risk [%]	Sharpe ratio	Turnover	Size of short side
	$N_{\text{estim}} = 75$				
SparseConc	14.376 (4.12)	19.035 (0.60)	**0.493** (0.22)	**0.186** (0.04)	**0.000** (0.00)
CondReg	13.256 (3.83)	17.681 (0.54)	0.467 (0.22)	0.593 (0.44)	0.043 (0.07)
Factor	9.581 (3.58)	15.722 (0.44)	0.291 (0.23)	1.809 (0.45)	0.219 (0.05)
LedoitWolf	9.687 (3.61)	15.350 (0.40)	0.305 (0.24)	1.785 (0.54)	0.201 (0.07)
Sample	9.545 (4.12)	17.486 (0.34)	0.260 (0.24)	2.673 (0.76)	0.303 (0.06)
	$N_{\text{estim}} = 150$				
SparseConc	14.353 (4.12)	19.043 (0.60)	**0.491** (0.22)	**0.177** (0.04)	**0.000** (0.00)
CondReg	12.917 (3.61)	16.493 (0.50)	0.480 (0.22)	0.921 (0.37)	0.090 (0.06)
Factor	9.690 (3.51)	15.397 (0.42)	0.305 (0.23)	1.781 (0.45)	0.213 (0.06)
LedoitWolf	9.837 (3.41)	14.868 (0.41)	0.325 (0.23)	1.748 (0.53)	0.194 (0.07)
Sample	9.487 (3.59)	15.728 (0.44)	0.285 (0.23)	2.144 (0.68)	0.254 (0.06)
	$N_{\text{estim}} = 225$				
SparseConc	14.407 (4.12)	19.053 (0.60)	**0.494** (0.22)	**0.173** (0.04)	**0.000** (0.00)
CondReg	12.934 (3.55)	16.131 (0.49)	0.492 (0.22)	1.035 (0.20)	0.102 (0.04)
Factor	10.011 (3.47)	15.256 (0.42)	0.328 (0.23)	1.751 (0.48)	0.208 (0.06)
LedoitWolf	10.411 (3.34)	14.690 (0.41)	0.368 (0.23)	1.719 (0.54)	0.190 (0.07)
Sample	10.482 (3.44)	15.144 (0.41)	0.362 (0.23)	1.980 (0.65)	0.235 (0.07)
	$N_{\text{estim}} = 225$				
SparseConc	14.403 (4.13)	19.085 (0.60)	0.493 (0.22)	**0.170** (0.03)	**0.000** (0.00)
CondReg	13.077 (3.54)	16.164 (0.49)	**0.499** (0.22)	1.037 (0.16)	0.098 (0.03)
Factor	10.656 (3.47)	15.262 (0.41)	0.371 (0.23)	1.735 (0.48)	0.205 (0.06)
LedoitWolf	10.680 (3.33)	14.660 (0.40)	0.387 (0.23)	1.699 (0.54)	0.186 (0.07)
Sample	10.809 (3.43)	15.076 (0.40)	0.385 (0.23)	1.890 (0.62)	0.223 (0.07)

6 Conclusion

This paper proposes a novel sparse inverse covariance estimation method G-AMA that adds and extends the sparse inverse covariance estimation literature in different directions. In particular, G-AMA is numerically faster than the previous best algorithms by several orders of magnitude in several regimes, while retaining all the good theoretical convergence properties of recently proposed methods. Moreover, G-AMA can solve ill-conditioned problems more accurately where existing methods do not converge or are extremely slow. The extra speed and ability to solve poorly conditioned problems makes G-AMA highly attractive for modern day applications, where the inverse covariance matrix has to be estimated multiple times. The G-AMA formulation also facilitates incorporating domain specific information while still maintaining linear convergence.

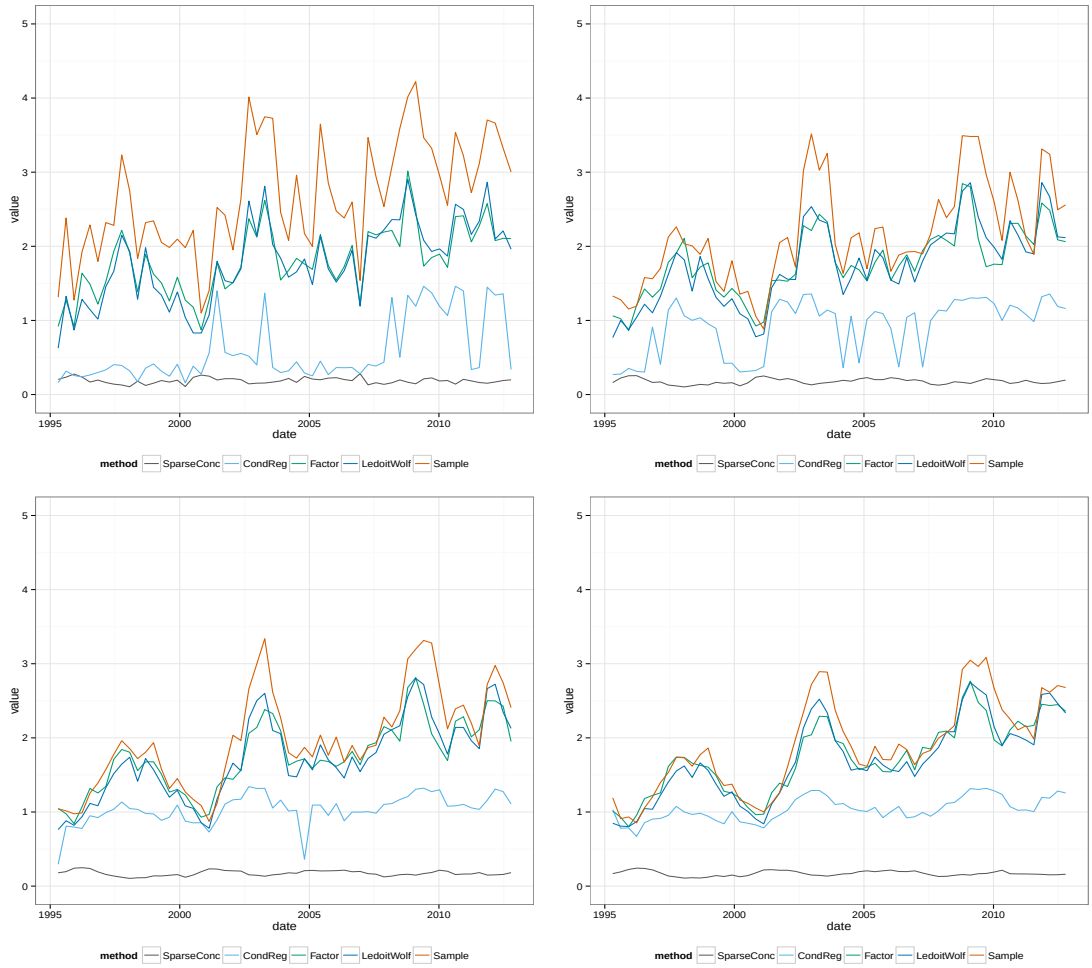

Figure 3: Turnover of G-AMA vs. other methods ($\lambda = 0.1 \times \|S\|_2$)

A Algorithm Details

A.1 Termination criteria

The algorithm is terminated when Z_k is positive definite and a tolerance of ϵ_{opt} is attained by the duality gap. The duality gap is given by

$$\Delta_{\text{opt}} = -\log\det Y_k - p - \log\det Z_k + \langle S, Z_k\rangle + \lambda\|Z_k\|_1.$$

When the algorithm terminates, the covariance estimate and the sparse inverse covariance estimate are given by Y_{k+1} and Z_{k+1} respectively.

An additional criteria of tolerance of ϵ_{prim} can be imposed on the progress of Y_k iterates

$$\Delta_{\text{prim}} = \frac{\|Y_{k+1} - Y_k\|_F}{\|Y_k\|_F}.$$

This quantity is related to the primal constraint violation as

$$\text{primal feasibility} = \|X_k - Z_k\|_F = \frac{\|Y_{k+1} - Y_k\|_F}{\tau}.$$

Therefore, the tolerance of ϵ_{prim} indirectly imposes primal feasibility.

A.2 Choice of step size

Recall that the step size has to be chosen in equation (5). The step size for each iteration is chosen starting with an adaptive Barzilia-Borwein (BB) step [2].

For a gradient descent algorithm minimizing a convex function $f(x)$ with gradient $\nabla f(x) = g$, the BB step size τ maximizes

$$\tau = \underset{\tau}{\text{argmax}} \|\tau \Delta x - \Delta g\|_2 = \frac{\langle \Delta x, \Delta x \rangle}{\langle \Delta g, \Delta x \rangle}, \tag{28}$$

which is a two-point approximation to the secant equation.

An equivalent step size for G-AMA solving 28 gives

$$\tau_{k+1} = \frac{\langle Y_{k+1} - Y_k, Y_{k+1} - Y_k \rangle}{\langle Y_{k+1} - Y_k, X_k - X_{k+1} \rangle} \tag{29}$$

is used in the numerical experiments. This heuristic provides substantial acceleration for convergence as compared to the fixed step size. However, the BB step might not always be feasible or satisfy sufficient descent condition. Hence, starting with the BB step we backtrack until the sufficient descent condition is satisfied.

B Lemmas and Proofs from Section 3

B.1 Proof of Lemma 3.1

Proof. Starting with X_* satisfying (8), we have three cases.
If $X_*(i,j) = 0$,

$$X_*^{-1}(i,j) - S(i,j) = \mathcal{C}_\lambda\left(\tau X_*(i,j) + (X_*^{-1}(i,j) - S(i,j))\right), \ \because \ \tau X_*(i,j) = 0.$$

Else if $X_*(i,j) > 0$ and $X_*^{-1}(i,j) - S(i,j) = \lambda$, we have

$$\mathcal{C}_\lambda\left(\tau X_*(i,j) + (X_*^{-1}(i,j) - S(i,j))\right) = \lambda, \ \because \ \tau X_*(i,j) > 0.$$

And if $X_*(i,j) < 0$ and $X_*^{-1}(i,j) - S(i,j) = -\lambda$, we have

$$\mathcal{C}_\lambda\left(\tau X_*(i,j) + (X_*^{-1}(i,j) - S(i,j))\right) = -\lambda, \ \because \ \tau X_*(i,j) < 0.$$

Hence, the inverse of X_* is a fixed point of (9).

For the reverse result, we begin with the inverse of a fixed point $X_* = Y_*^{-1}$. Therefore, we have,

$$X_*^{-1} - S = \mathcal{C}_\lambda\left(\tau X_* + (X_*^{-1} - S)\right)$$

The entries on the right hand side are contained in the interval $[-\lambda, \lambda]$ and the additional term of τX_* in the argument of clip function implies that $X_*(i,j)$ has to be 0 when $|X_*^{-1} - S| < \lambda$ and $X_*(i,j)$ is non-zero only if $(X_*^{-1} - S)(i,j) = \pm\lambda$ with same sign as $X_*(i,j)$. Thus $X_*^{-1} - S$ satisfy both the conditions in (8) proving the required result. □

B.2 Contraction of G-AMA iterates

Using Lemma 3.2, we now prove the following lemma which is a key ingredient for showing upper bound on iterates of G-AMA as well as global linear convergence of G-AMA.

Lemma B.1. *Let*

$$Y_+ = \mathcal{C}_\lambda\left(Y + \tau Y^{-1} - S\right) + S, \tag{30}$$

be from (5) and let Y_* be a fixed point from equation (9), i.e.,

$$Y_* = \mathcal{C}_\lambda \left(Y_* + \tau Y_*^{-1} - S \right) + S. \tag{31}$$

Also, define

$$\alpha = \min \left\{ \lambda_{\min}(Y), \lambda_{\min}(Y_*) \right\}, \beta = \max \left\{ \lambda_{\max}(Y), \lambda_{\max}(Y_*) \right\}.$$

Then

$$\|Y_+ - Y_*\|_F \leq \max \left\{ \left| 1 - \frac{\tau}{\alpha^2} \right|, \left| 1 - \frac{\tau}{\beta^2} \right| \right\} \|Y - Y_*\|_F.$$

Proof. We begin with the left hand side and substituting (5) and (9) and using the non-expansive property (see Lemma 4.1) of the $\mathcal{C}_\lambda$ function (projection onto a convex set), we get

$$\begin{aligned} \|Y_+ - Y_*\|_F &= \|\mathcal{C}_\lambda \left(Y + \tau Y^{-1} - S \right) - \mathcal{C}_\lambda \left(Y_* + \tau Y_*^{-1} - S \right) \|_F, \\ &\leq \| \left(Y + \tau Y^{-1} \right) - \left(Y_* + \tau Y_*^{-1} \right) \|_F, \end{aligned} \tag{32}$$

To bound the later expression we follow the arguments from [9, Lemma 3] and define a function

$$h_\tau(Y) = \text{vec}(Y) + \text{vec}(\tau Y^{-1}).$$

Using h_τ, (32) can be written as

$$\| \left(Y + \tau Y^{-1} \right) - \left(Y_* + \tau Y_*^{-1} \right) \|_F = \|h_\tau(Y) - h_\tau(Y_*)\|_2.$$

The Jacobian J_{h_τ} of h_τ is given by

$$J_{h_\tau}(Y) = I_{p^2} - \tau Y^{-1} \otimes Y^{-1}.$$

The function $h_\tau(Y)$ is a differentiable on $\mathcal{D}$. Therefore applying the mean value theorem gives

$$\begin{aligned} \|h_\tau(Y) - h_\tau(Y_*)\|_2 &\leq \sup_{\delta \in [0,1]} \left\{ \|I_{p^2} - \tau Y_\delta^{-1} \otimes Y_\delta^{-1}\|_2 \right\} \|\text{vec}(Y) - \text{vec}(Y_*)\|_2, \\ &\leq \sup_{\delta \in [0,1]} \left\{ \|I_{p^2} - \tau Y_\delta^{-1} \otimes Y_\delta^{-1}\|_2 \right\} \|Y - Y_*\|_F, \end{aligned}$$

where $Y_\delta = \delta Y + (1 - \delta) Y_*$ is some convex combination of Y and Y_*. The eigenvalues of Y_δ are

bounded using Weyl's inequality which gives

$$\alpha I \preceq Y_\delta \preceq \beta I.$$

The eigenvalues of $I_{p^2} - \tau Y_\delta^{-1} \otimes Y_\delta^{-1}$ are $1 - \tau\left[\text{eig}(Y_\delta)^{-2}\right]$ and therefore the $\sup$ term can be bounded as

$$\sup_{\delta \in [0,1]} \left\{ \| I_{p^2} - \tau Y_\delta^{-1} \otimes Y_\delta^{-1} \|_2 \right\} = \max \left\{ \left| 1 - \frac{\tau}{\alpha^2} \right|, \left| 1 - \frac{\tau}{\beta^2} \right| \right\},$$

thereby completing the proof. □

B.3 Proof of Lemma 3.3

Proof. We begin with the maximum eigenvalue of Y_{k+1}.

$$\begin{aligned} \lambda_{\max}(Y_{k+1}) = \|Y_{k+1}\|_2 &= (\|Y_{k+1}\|_2 - \|Y_*\|_2) + \|Y_*\|_2 \\ &\leq \|Y_{k+1} - Y_*\|_2 + \|Y_*\|_2 \\ &\leq \|Y_{k+1} - Y_*\|_F + \|Y_*\|_2 \end{aligned}$$

Since $\tau_k < \alpha_k^2 \leq \beta^2$,

$$\max \left\{ \left| 1 - \frac{\tau_k}{\alpha_k^2} \right|, \left| 1 - \frac{\tau_k}{\beta^2} \right| \right\} \leq 1,$$

and therefore,

$$\lambda_{\max}(Y_{k+1}) \leq \|Y_k - Y_*\|_F + \|Y_*\|_2,$$

and applying Lemma B.1 repeatedly for $\|Y_k - Y_*\|_F \leq \|Y_{k-1} - Y_*\|_F$, we have

$$\lambda_{\max}(Y_{k+1}) \leq \|Y_0 - Y_*\|_F + \|Y_*\|_2,$$

which proves the required result with inductive argument on k. □

B.4 Proof of Lemma 3.4

Proof. Denote by $a = \lambda_{\min}(Y)$ the smallest eigenvalue of Y. By Lemma 3.3 we have

$$\log\det Y_0 < \log\det Y \leq \log(a) + (p-1)\log(\beta)$$

For the initial point $Y_0 = S + \lambda I$, the minimum eigenvalue of Y_0 satisfies $\lambda \leq \lambda_{\min}(Y_0)$ and $p\lambda \leq \log\det Y_0$. Therefore,

$$\lambda_{\min}(Y) = a > (\det Y_0)^p \, \beta^{-(p-1)} \geq \lambda^p \beta^{-(p-1)}.$$

Hence proved. □

B.5 Proof of Theorem 3.5

Proof. The iterations in Algorithm 1 can be rewritten by substituting X_{k+1} and Z_{k+1} in terms of Y_k as

$$\begin{aligned} Y_{k+1} &= Y_k + \tau_k \left(Y_k^{-1} - \mathcal{S}_{\lambda/\tau}(Y_k^{-1} + (Y_k - S)/\tau_k) \right), \\ &= \mathcal{C}_\lambda \left(Y_k + \tau Y_k^{-1} - S \right) + S. \end{aligned}$$

The initial point

$$Y_0 = S + \lambda I,$$

and the subsequent iterates satisfy $\alpha I \preceq Y_k \preceq \beta I$ by Lemma 3.3 and Lemma 3.4 for α, β defined in the corresponding lemmas. Therefore, using Lemma B.1, we get equation (26)

$$\|Y_{k+1} - Y_*\|_F \leq \max\left\{ \left| 1 - \frac{\tau_k}{\alpha^2} \right|, \left| 1 - \frac{\tau_k}{\beta^2} \right| \right\} \|Y_k - Y_*\|_F,$$

for $k = 0, 1 \ldots$. And the condition on the constant step size $\tau_k < \alpha^2, \forall k$ gives (14),

$$\|Y_{k+1} - Y_*\|_F \leq \gamma \|Y_k - Y_*\|_F,$$

thereby proving the linear convergence of Algorithm 1. □

B.6 Proof of Theorem 4.2

Proof. (a) Using Lemma 4.1, the modified Y_+ update (18) satisfies

$$\begin{aligned}\|Y_+ - Y_*\|_F &= \|\Pi_{\mathcal{D}}(Y + \tau Y^{-1}) - \Pi_{\mathcal{D}}(Y_* + \tau Y_*^{-1})\|_F \\ &\leq \| \left(Y + \tau Y^{-1}\right) - \left(Y_* + \tau Y_*^{-1}\right) \|_F, \end{aligned} \tag{33}$$

and the rest of the proof of Lemma B.1 follows accordingly. Similarly the proof of Theorem 3.5 can easily be adapted to prove the linear convergence of the modified algorithm.
(b) To prove the convergence of (23), we use a similar mean-value theorem argument on the function $Y + \tau A^*(S + BYA^*)^{-1}B$, namely,

$$\begin{aligned}\|Y_+ - Y_*\|_F &\leq \|Y - Y_* + \tau A^* \left((S + BYA^*)^{-1} - (S + BY_*A^*)^{-1}\right) B\|_F \\ &\leq \max\left\{|1 - \tau v|, |1 - \tau w|\right\} \|Y - Y_*\|_F, \end{aligned} \tag{34}$$

where v and w are constants which depend on S, $\rho(A^*A)$, $\rho(B^*B)$ and λ. □

Appendix C: Portfolio Optimization Background

The formal definitions of the performance metrics used in Section 5.3 as taken from [27] are as follows

- *Realized return.* The realized return of a portfolio rebalancing strategy over the entire trading period is

$$r_p = \frac{1}{K}\sum_{j=1}^{K}\frac{1}{L}\sum_{t=N_{\text{estim}}+1+(j-1)L}^{N_{\text{estim}}+jL} r^{(t)T}w^{(j)}.$$

- *Realized risk.* The realized risk (return standard deviation) of a portfolio rebalancing strategy over the entire trading period is

$$\sigma_p = \sqrt{\frac{1}{K}\sum_{j=1}^{K}\frac{1}{L}\sum_{t=N_{\text{estim}}+1+(j-1)L}^{N_{\text{estim}}+jL} \left(r^{(t)T}w^{(j)}\right)^2 - r_p^2}.$$

- *Realized Sharpe ratio (SR).* The realized Sharpe ratio, *i.e.*, the ratio of the excess expected return of a portfolio rebalancing strategy relative to the risk-free return r_f is given by

$$SR = \frac{r_p - r_f}{\sigma_p}.$$

- *Turnover.* The turnover from the portfolio $w^{(j)}$ held at the start date of the jth holding period $[N_{\text{estim}} + 1 + (j-1)L, N_{\text{estim}} + jL]$ to the portfolio $w^{(j-1)}$ held at the previous period is computed as

$$\text{TO}(j) = \sum_{i=1}^{p} \left| w_i^{(j)} - \left(\prod_{t=N_{\text{estim}}+1+(j-1)L}^{N_{\text{estim}}+jL} (1 + r_i^{(t)}) \right) w_i^{(j-1)} \right|.$$

  For the first period, we take $w^{(0)} = 0$, *i.e.*, the initial holdings of the assets are zero.

- *Normalized wealth growth.* Let $w^{(j)} = (w_1^{(j)}, \ldots, w_n^{(j)})$ be the portfolio constructed by a rebalancing strategy held over the jth holding period $[N_{\text{estim}} + 1 + (j-1)L, N_{\text{estim}} + jL]$. When the initial budget is normalized to one, the normalized wealth grows according to the recursion

$$W(t) = \begin{cases} W(t-1)(1 + \sum_{i=1}^{p} w_{it} r_i^{(t)}), & t \notin \{N_{\text{estim}} + jL \mid j = 1, \ldots, K\}, \\ W(t-1)\left((1 + \sum_{i=1}^{p} w_{it} r_i^{(t)}) - \text{TC}(j)\right), & t = N_{\text{estim}} + jL, \end{cases}$$

  for $t = N_{\text{estim}}, \ldots, N_{\text{estim}} + KL$, with the initial wealth $W(N_{\text{estim}}) = 1$. Here

$$w_{it} = \begin{cases} w_i^{(1)}, & t = N_{\text{estim}} + 1, \ldots, N_{\text{estim}} + L, \\ \vdots & \\ w_i^{(K)}, & t = N_{\text{estim}} + 1 + (K-1)L, \ldots, N_{\text{estim}} + KL. \end{cases}$$

  and

$$\text{TC}(j) = \sum_{i=1}^{p} \eta_i \left| w_i^{(j)} - \left(\prod_{t=N_{\text{estim}}+1+(j-1)L}^{N_{\text{estim}}+jL} (1 + r_i^{(t)}) \right) w_i^{(j-1)} \right|$$

  is the transaction cost due to the rebalancing if the cost to buy or sell one share of stock i is η_i.

- *Size of the short side.* The size of the short side of a portfolio rebalancing strategy over the entire trading period is computed as

$$r_p = \frac{1}{K} \sum_{j=1}^{K} \left(\sum_{i=1}^{p} |\min(w_i^{(j)}, 0)| / \sum_{i=1}^{p} |w_i^{(j)}| \right).$$

Acknowledgment

We would like to thank Dominique Guillot for useful discussions and suggestions for improving the paper. We would also like to thank the authors of [9] and [11] for providing the code for G-ISTA and QUIC which served as a base for the code development for G-AMA.

References

[1] O. Banerjee, L. El Ghaoui, and A. d'Aspremont. Model selection through sparse maximum likelihood estimation for multivariate gaussian or binary data. JMLR, 9:485–516, 2008.

[2] J. Barzilai and J.M. Borwein. Two-point step size gradient methods. IMA Journal of Numerical Analysis, 8(1):141–148, 1988.

[3] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. Foundations and Trends® in Machine Learning, 3(1):1–122, 2011.

[4] S. Boyd and L. Vandenberghe. Convex optimization. Cambridge university press, 2004.

[5] P. Brohan, J.J. Kennedy, I. Harris, S.F.B. Tett, and P.D. Jones. Uncertainty estimates in regional and global observed temperature changes: A new data set from 1850. Journal of Geophysical Research, 111(D12):D12106, 2006.

[6] A. d'Aspremont, O. Banerjee, and L. El Ghaoui. First-order methods for sparse covariance selection. SIAM Journal on Matrix Analysis and Applications, 30(1):56–66, 2008.

[7] J. Friedman, T. Hastie, and R. Tibshirani. Sparse inverse covariance estimation with the graphical lasso. Biostatistics, 9(3):432, 2008.

[8] Helene Gehrmann and Steffen L Lauritzen. Estimation of means in graphical gaussian models with symmetries. The Annals of Statistics, 40(2):1061–1073, 2012.

[9] D. Guillot, B. Rajaratnam, B. Rolfs, , A. Maleki, and I. Wong. Iterative thresholding algorithm for sparse inverse covariance estimation. Advances in Neural Information Processing Systems, 2012.

[10] Cho-Jui Hsieh, Inderjit Dhillon, Pradeep Ravikumar, and Arindam Banerjee. A divide-and-conquer method for sparse inverse covariance estimation. In Advances in Neural Information Processing Systems 25, pages 2339–2347, 2012.

[11] C.J. Hsieh, M.A. Sustik, I.S. Dhillon, and P. Ravikumar. Sparse inverse covariance matrix estimation using quadratic approximation. Advances in Neural Information Processing Systems (NIPS), 24, 2011.

[12] Kshitij Khare, Sang-Yun Oh, and Bala Rajaratnam. A convex pseudo-likelihood framework for high dimensional partial correlation estimation with convergence guarantees. arXiv preprint arXiv:1307.5381, 2013.

[13] O. Ledoit and M. Wolf. Honey, i shrunk the sample covariance matrix. UPF Economics and Business Working Paper, (691), 2003.

[14] O. Ledoit and M. Wolf. A well-conditioned estimator for large-dimensional covariance matrices. Journal of multivariate analysis, 88(2):365–411, 2004.

[15] O. Ledoit and M. Wolf. Robust performance hypothesis testing with the sharpe ratio. Journal of Empirical Finance, 15(5):850–859, 2008.

[16] J.D. Lee, Y. Sun, and M.A. Saunders. Proximal newton-type methods for minimizing convex objective functions in composite form. arXiv preprint arXiv:1206.1623, 2012.

[17] Z. Lu. Smooth optimization approach for sparse covariance selection. SIAM Journal on Optimization, 19(4):1807–1827, 2009.

[18] Z. Lu. Adaptive first-order methods for general sparse inverse covariance selection. SIAM Journal on Matrix Analysis and Applications, 31(4):2000–2016, 2010.

[19] R. Mazumder and T. Hastie. The graphical lasso: New insights and alternatives. Electronic Journal of Statistics, 6:2125–2149, 2012.

[20] S. Oh, B. Rajaratnam, and J. Won. Condreg: An r package for condition number regularized covariance estimation. 2012.

[21] J. Pittman, E. Huang, H. Dressman, C.F. Horng, S.H. Cheng, M.H. Tsou, C.M. Chen, A. Bild, E.S. Iversen, A.T. Huang, et al. Integrated modeling of clinical and gene expression information for personalized prediction of disease outcomes. Proceedings of the National Academy of Sciences of the United States of America, 101(22):8431–8436, 2004.

[22] R.T. Rockafellar. Monotone operators and the proximal point algorithm. SIAM Journal on Control and Optimization, 14(5):877–898, 1976.

[23] R.T. Rockafellar and R.J.B. Wets. Variational analysis, volume 317. Springer, 2011.

[24] Katya Scheinberg, Shiqian Ma, and Donald Goldfarb. Sparse inverse covariance selection via alternating linearization methods. arXiv preprint arXiv:1011.0097, 2010.

[25] Tapio Schneider. Analysis of incomplete climate data: Estimation of mean values and covariance matrices and imputation of missing values. Journal of Climate, 14(5):853–871, 2001.

[26] P. Tseng. Applications of a splitting algorithm to decomposition in convex programming and variational inequalities. SIAM Journal on Control and Optimization, 29(1):119–138, 1991.

[27] J. Won, S.J. Kim, J. Lim, and B. Rajaratnam. Condition number-regularized covariance estimation. Journal of the Royal Statistical Society: Series B (Statistical Methodology), 75(3):427–450, 2013.